

The Economist

10-16 May 1986

SCIENCE AND TECHNOLOGY

The Connection Machine starts to think in parallel


If you have \$3m to spare and a penchant for dashing off designs for integrated circuits, searching databases the size of municipal libraries or modelling images in three dimensions, you should invest in a black, five-foot-high cube called the Connection Machine. Made by a small Boston company called Thinking Machines, it was launched last week some years after the first rumours of its power began circulating. One breathless customer, Perkin-Elmer (a computer maker in Northwalk, Connecticut), calls it "the most significant single major advance in large-scale computing in several decades".

The Connection Machine consists of 65,536 minute computers (ie, microprocessors with attached memories) all wired together. Conventional computers—using the so-called von Neumann architecture—have one processor fed by a large memory. The data required for each bit of a computation are shipped into the processor, which does what is needed and ships them back out again. This means that the processor's speed and—more crucially—the speed of the links put an upper limit on the processing power of such a machine. That limit has almost been reached.

There have been several attempts to break this bottleneck and build a "non-von" machine. One is to have several processors, each working on different bits of the problem. Many expensive machines now incorporate more than one processor. Even Crays, archetypal von Neumann machines which fill rooms, cost tens of millions of dollars and crunch numbers at prodigious rates, now have several. Companies like Concurrent (a Perkin-Elmer offshoot) and the Oregon-based Sequent specialise in adding performance to cheaper machines in this way.

The Connection Machine takes multiprocessing to its logical extreme: it splits

the problem into thousands of parts and deals with them in parallel. One reason for thinking this is a good idea is that human brains do the same. Mr Danny Hillis sketched out the essentials of such a machine in a doctoral thesis at the Massachusetts Institute of Technology. His ideas were made silicon through the efforts of Dr Sheryl Handler, who put together an impressive team of experts in artificial intelligence, computation and circuit design, and then found \$16m from



Connections as far as the eye can see

venture capitalists for the new company.

Each processor in the machine is tiny: it can handle "words" of only one binary digit (bit) at a time (the new generation of processors handles 32-bit words), and stores only about 4,000 bits of information on a chip. However, this is enough for all sorts of applications, and such small units come cheap. Many kinds of big problem can be analysed by looking at thousands of tiny sub-problems. A good example is the computer modelling of fluid flows, such as air passing over an

aircraft wing. It is relatively easy to calculate how each bit of air is affected by the forces acting on it. It takes millions of calculations—but each of them is simple. Von Neumann machines are poorly qualified to deal with this sort of problem efficiently: most of the expensive beast stands idle while its brain works at a furious rate. The Connection Machine was designed with such problems in mind.

The machine has a number of features that make it more sophisticated than previous attempts at parallelism. Most important is its flexibility. People have tried linking processors in rings, lines, grids and hypercubes (cubes which manage to have more than three dimensions). Mr Hillis stresses that he has made all the links in the Connection Machine programmable with software: the machine

hooks itself up, along the lines of the packet switching system used in telecommunications, in whichever way is best suited to the problem in hand. Thinking Machines calls this "data-level parallelism", because the links follow the size and shape of the data.

A big criticism of parallel machines is that they will require new software. The Connection Machine proves that this is nonsense: it runs in the familiar programming languages LISP and C. Indeed, it is designed as a sort of fast processing extension to an ordinary mini-computer, like a DEC VAX or a Symbolics LISP machine. Turning a conventional sequential program into something which will run in parallel is done by the software rather than the user. The user is aware only that his mini has magically speeded up.

The Connection Machine is far from being 65,000 times faster than a conventional machine. But its performance is impressive, especially considering its \$3m price tag. Mr Hillis reckons that his machine is up to twice as fast at modelling fluid dynamics as the latest Cray—which costs about \$10m. And the Cray is tailored to this kind of problem, whereas the Connection Machine, because of its flexible structure, is a polymath.

In one test, which involved searching a database of 10 billion characters (or about

Reprinted from THE WALL STREET JOURNAL.

THURSDAY, MAY 1, 1986

© 1986 Dow Jones & Company, Inc. All Rights Reserved.

Thinking Machines Unveils a Computer With New Technique

By a WALL STREET JOURNAL Staff Reporter

CAMBRIDGE, Mass. — Thinking Machines Corp. unveiled its first product, a computer called the Connection Machine that uses a technique called parallel processing to solve some problems much faster than traditional computers.

The closely held, three-year-old company has attracted widespread attention in the computer world because of reports of the speed and unusual design of its computer. It was the subject of a front-page article in The Wall Street Journal in February.

Thinking Machines announced two models of its computer, one with 65,536 processors, which will sell for \$3 million, and one with 16,384 processors, which will cost \$1 million. The individual processors work together on problems such as modeling the flow of water molecules past a boat hull or designing the circuits in a semiconductor. Although each processor is small, in combination they can process one billion instructions per second, the company said. The biggest conventional computers process about 40 million instructions per second.

The company, backed by \$16 million in venture capital, grew out of work at the Massachusetts Institute of Technology. It said early customers for the machine, which will be available in July, include two MIT laboratories, the Defense Department, Yale University, and Perkin-Elmer Corp., which will use one to study computer vision, artificial intelligence and simulation.

Scientists have worried that parallel processing would prove unworkable because programming would be so difficult. However, Oliver McBryan, a computer scientist with New York University who has worked with the machine, said it "is an order of magnitude easier to program," than previous parallel processors.

nature

NATURE VOL. 321 8 MAY 1986

US technology

Parallel computer unveiled

Cambridge, Massachusetts

Just two and a half years after setting out to build the world's first "fifth generation" computer, Thinking Machines Corporation (TMC) of Cambridge, Massachusetts, last week unveiled the Connection Machine (CM), a massively parallel computer with 65,536 individual processors. TMC has already developed some remarkable applications for the new machine but the company consultant, Richard Feynman of California Institute of Technology, says the problems have yet to be dreamed up that will make full use of CM's unique abilities.

While a graduate student at Massachusetts Institute of Technology (MIT) in Cambridge, CM's spiritual father Danny Hillis came up with the blueprint for a data-level parallel processing computer. Using data-level processing leads to a more natural way of thinking about certain problems, says Hillis. For example, organizing a visual image by examining one small section at a time makes less sense than processing all the information in the image simultaneously. Data-level processing works best on problems with large amounts of data, whereas control-level parallel processing machines are most effective when the ratio of program to data is high. But control-level parallelism involves the difficult task of developing algorithms for asynchronous control of operations.

Communication among processors is the key to CM. Locally, this is done by interconnections among neighbouring processors. For general communications and dynamic reconfiguration, CM uses routers to send messages to the 64K pro-

cessors arranged in a 16-dimensional hypercube at speeds of up to 3,000 megabits per second.

CM consists of a front-end computer, either a VAX or a Symbolics 3600, that sends instructions to the parallel processors. Each processor contains 4K of memory, so a fully configured CM has 32 megabytes of memory. Even with a \$3 million price tag, or \$1 million for a CM with 16K processors, prospective buyers are lining up to get hold of the new machine. TMC has so far delivered only one CM, to the Defense Advanced Research Projects Agency (DARPA), but in two months machines will be heading to MIT (2), Yale, Perkin Elmer Corporation and a second machine to DARPA.

In 1983, Hillis teamed up with Sheryl Handler to form TMC. Handler wanted to find a product that would provide a "fundamentally new way of computing". The new company quickly brought on board top scientists as consultants: Hillis's adviser Marvin Minsky and Thomas Poggio of MIT, Stephen Wolfram of the Center for Advanced Study at Princeton and Caltech's Feynman. With \$16 million in equity capital and a \$4.7 million DARPA contract early on, TMC was able to take the risks Handler felt necessary to leapfrog the competition.

A major advantage of CM for new users is its ability to run in familiar operating environments. Extensions to C and LISP computer languages designed by TMC allow immediate access to CM's parallel capabilities. TMC has already come up with applications packages that show off CM's power. A search for key

words in a 16,000-document database takes only 30 milliseconds, much faster than conventional searches. Image processing, very large scale integrated (VLSI) circuit design, and fluid flow problems all take advantage of CM by assigning one processor to each element of the problem's dataset. CM can achieve operating speeds up to 7,000 MIPS (million instructions per second) on certain applications, faster than the fastest computer utilizing more conventional architecture.

Although CM's initial capabilities are impressive, Wolfram believes that a new computer language is needed to take full advantage of it. CM is well suited to cellular automata, and Wolfram predicts that almost all problems now being solved using differential equations will be done in future by constructs like cellular automata.

Another plus for CM is its expandability. While traditional computers using principles developed 40 years ago by von Neumann are inherently limited by a single central processing unit, an arbitrarily large number of processors can in theory be linked together using CM's design. TMC scientists have already written simulations of a machine with one million processors that can run on current CM computers.

Hillis says TMC's next goal is to develop a true learning system. In a small way, the interconnecting processors of the CM resemble the interconnections of individual neurones of the brain. Hillis hopes that this type of architecture will encourage the development of learning algorithms that more closely resemble biological processes.

Other groups around the world are working on new computers using parallel architecture, but Poggio points out that CM has one obvious advantage over other projects: "It works". **Joseph Palca**

Reprinted with permission.

Thinking Machines Corporation

245 First Street, Cambridge, MA 02142-1214 617 • 876 • 1111

Computer firm banks on speed

By Ronald Rosenberg
Globe Staff

Thinking Machines Corp. put some sparkle in the computer industry yesterday with a super-fast computer that it says will change the way traditional computer functions are done.


At a press conference in the company's Cambridge headquarters along the banks of the Charles River, company officials led by Sheryl L. Handler, the founder and president, introduced the Connection Machine. It is a computer that can process information at the rate of 1 billion instructions per second. (By comparison, International Business Machines Corp.'s top-of-the-line Sierra mainframe is rated at 28 to 40 million instructions per second).

The high speeds can reduce the search and retrieval of a voluminous database to seconds instead of many minutes. To electrical engineers the Connection Machine promises to speed up the simulation of electronic circuits. Scientists can do complex modeling used in weather forecasting, automobile streamlining and airplane design more expeditiously.

"This is quite an achievement," said Jacob T. Schwartz, a professor of mathematics and computer science at New York University who developed some of the Connection Machines' algorithms as a consultant.

Distinguishing the Connection Machine from other computers is the interconnection of between 16,000 and 64,000 custom-designed microprocessors. Previous machines have used no more than several hundred processors. Unlike conventional computers that process information in a sequence, piece by piece, this computer processes all the data at the same time. The high speed is achieved by connecting the microprocessors, which act on the information and communicate with other microprocessors.

The five-foot-high computer is an imposing machine. It consists of eight black cubes, each with 512 flashing red lights that are visible through a semitranspar-



Thinking Machines Corp. new Connection Machine system.

ent surface at one end. Although the lights aid in computer maintenance they are used largely for marketing.

"I love the lights," chuckled W. Daniel Hillis, the 29-year-old architect of the Connection Machine who first explored the concept while a graduate student at Massachusetts Institute of Technology. "I want to know when the computer is working, and looking at the video display terminal is not enough."

His doctoral thesis adviser was Marvin Minsky, considered the father of the artificial-intelligence field, who like Hillis is a founding scientist of the company. Minsky predicted similar machines will be on desktops in 10 years.

Both men, together with Sheryl L. Handler, who founded Thinking Machines in June 1983, expected the research and development to take about five years. She, too, is an MIT graduate and a businesswoman.

"Where we are today is what I expected would take five years," said Handler, who drew some criticism by attracting scientists to leave academia and work for Thinking Machines. "We couldn't

have accomplished the Connection Machine if we relied on university research. We needed to bring the best people together."

Thinking Machines also attracted \$16 million in capital from such investors as William Paley, the founder of CBS who is also an investor in another Cambridge company, Genetics Institute.

"Having the top scientists impressed me," said Paley, who confers with Handler at least once a month. "Scientists are pretty careful and quiet guys. This field is a lot different from the business and entertainment field."

Critics of the Connection Machine claim its high price of \$1 million to \$3 million and the need for specially adapted software may limit its impact. Others point to the wave of minisupercomputer machines coming this year that although less powerful are also less costly.

But others contend it will be a winner.

The company has sold six machines, including one each to MIT and Yale University.

Reprinted courtesy of The Boston Globe.

The Boston Herald

Thursday May 1, 1986

New computer closest thing to human brain

By GEOFFREY ROWAN

A COMPUTER which can solve problems in seconds that take traditional computers hours to unravel was unveiled yesterday by a Cambridge company.

The Connection Machine, by Thinking Machines Inc., uses 64,000 separate processors in unison to create a computer which thinks more like a human than a standard machine.

Comparing the speed of traditional computers to the Connection Machine is like "comparing a bicycle to a supersonic jet," said the computer's designer, W. Daniel Hillis.

Its incredible power — the Connection Machine can process 1 billion instructions per second — comes from a technology called parallel processing.

Conventional computers use a single, large processor to compute information in a linear fashion, extract-

ing one piece of data from its memory, processing that and returning to its memory before moving on to the next part of the problem.

The Connection Machine uses 64,000 processors working on different pieces of a problem at the same time, much as the human brain depends upon billions of neurons working at the same time to keep many different human functions going.

"The result is a computer that looks at the whole problem at once," Hillis said.

Other companies have attempted to create parallel processing machines but they have worked on a smaller scale. Most are considered multiple processors because while they use several processors they do not operate in parallel.

Thinking Machines said it will market the \$3 million Connection Machine to gov-

ernment agencies, Fortune 500 companies with large research departments and universities.

The first six machines have been purchased by the Perkin-Elmer Corp., the Media Laboratory at the Massachusetts Institute of Technology, the Artificial Intelligence Laboratory at MIT, Yale University and the government's Defense Advanced Research Project Agency (DARPA), which has already received one and has another on order.

"We think the Connection Machine system is the wave of the future," said Prof. Patrick Winston, director of MIT's Artificial Intelligence Laboratory.

"The existence of a machine like this changes the way we think about problems," he said. "It will open up new ways of thinking in many areas."

Reprinted with permission.

Thinking Machines Corporation

245 First Street, Cambridge, MA 02142-1214 617 • 876 • 1111

A 1-BIPS SYSTEM TAKES NEW TACK IN PARALLELISM

CAMBRIDGE, MASS.

Mips worshipers have a new god. Thinking Machines Corp.'s Connection Machine, a massively parallel computer that made its commercial debut last week, goes beyond millions of instructions per second and ushers in the era of bips with its ability to execute 1 billion instructions/s.

Other numbers describing the Cambridge company's machine are likewise grand in scale. In its largest configuration, the computer has 65,536 processors and 32 megabytes of memory. Processors can communicate with each other at 3,000 Mb/s and the computer's input/output channel can handle 500 Mb/s. The air-cooled computer, housed in a cluster of eight translucent black cubes clustered one on one in four stacks with dozens of blinking red lights, needs a hefty front-end machine, either a Digital Equipment Corp. VAX or Symbolics Inc. computer.

More important than the raw numbers, however, is that the computer takes a new approach to parallelism. The Connection Machine uses data-level parallelism by reducing tasks to data elements and assigning a single processor to each element, says W. Daniel Hillis, founding scientist at Thinking Machines. This makes it particularly useful for applications involving huge amounts of data, as opposed to heavy computing on a few data elements.

As an example, Hillis describes an attempt to determine which of 50,000 news articles relates to a particular clipping. In the Connection Machine, he explains, each article would be assigned to and examined by a single processor. The processors would then compare their data elements with the reference clipping stored in the front-end computer.



HILLIS. The Connection Machine devotes a processor to each data element in a problem.

ALL AT ONCE. "The most rational way to execute a lot of operations is to work all the data at once," says Hillis. This fine-grain parallel-computing method also avoids the primary bugaboo of so-called coarse-grain processors (those with fewer, larger elements), which require subdivision of the application problem, he adds. "We don't partition, and that avoids a big set of problems."

The Connection Machine's cornerstone is a custom chip that holds 16 processors and 4-K of memory. If an application takes only some processors, the system temporarily switches off unused chips. If there are more data elements than processors, the computer's hardware operates in virtual-processor mode by subdividing memory and simulating additional multiple processors, each with a smaller memory. The Connection Machine can support up to a million virtual processors. The machine's breakthrough concept was eliminating the separation

of processors and memory and instead mixing them together along with high-speed communication elements, says president Sheryl Handler.

The Connection Machine runs under conventional operating systems, including AT&T Co.'s Unix, DEC's VMS, and Symbolics's Lisp environment. Its languages are extended versions of Lisp and C. Throughout development, says Handler, the development team concentrated on applications. At the large conference announcing availability of the machine, demonstration areas included document processing, contour mapping, chip design, and fluid dynamics.

The company also announced that several machines had already been sold, in addition to one delivered last fall to the Defense Advanced Research Projects Agency, which partly funded development of the computer and has placed an order for another machine. Other early customers are Perkin-Elmer Corp., the Massachusetts Institute of Technology, and Yale University.

Perkin-Elmer's 16,000-processor unit will be used at MRJ Inc., Oakton, Va., to search text, process images, and do research, says Edward McMahon, a member of the Perkin-Elmer division's technical staff. Referring to the difficulty of programming other multiprocessors, McMahon says, "The creativity in using the Connection Machine is in formulating algorithms, rather than the programming part."

Price for the computer with 65,536 processors and 32 megabytes of memory is \$3 million. A unit with 16,384 processors and 8 megabytes of memory sells for \$1 million. This compares favorably with large mainframes, claims vice president Richard Clayton. He calculates the Connection Machine's cost at \$3,000 per mips or less, whereas mainframes cost up to \$150,000/mips. To keep costs down, he explains, the company placed a priority on using simple building blocks and conservative manufacturing technology. —Craig D. Rose

Reprinted from **ELECTRONICS**, May 5, 1986, copyright 1986 by McGraw-Hill, Inc. with all rights reserved.
Additional reprints may be ordered by calling Electronics Reprint Department (609) 426-5494.

Thinking Machines Corporation
245 First Street
Cambridge, MA 02142-1214

Thinking Machines' Parallel Processor Connects Over 64,000 Nodes

By R. Colin Johnson

CAMBRIDGE, Mass. — Thinking Machines Corp. will launch its 64,000-node parallel processor here this week.

The success of the Connection Machine System in experimental work sponsored by the Defense Department led to Thinking Machines' release of a commercial version (see *EE Times*, July 29, Page 1).

The Connection Machine is a so-called fine granularity parallel processor, so called because it harnesses upwards of 64,000 nodes in its base configuration. Future versions are envisioned that could use a million or more nodes for supercomputer performance.

The Connection Machine currently has the finest granularity of any commercial parallel processor, though NCR has a new chip from which a system that rivals the Connection Machine in granularity could be built.

NCR's geometric arithmetic parallel processor, GAPP has 72 processors on a CMOS chip.

The Connection Machine System keeps its processors doing useful work with a VLSI interconnection scheme that mirrors the job being run. Nodes and their interconnections are closely matched to particular applications.

Artificial-intelligence researchers have been especially keen on the Connection Machine System as an expert system engine. "We are aiming at using the Connection Machine to retrieve elements of a knowledge base in just a few machine cycles rather than seconds or minutes," said Edward Feigenbaum, the author of *The Fifth Generation* and a DOD researcher.

The trick is that each node in

the Connection Machine System is set up to hold one element of knowledge. Then its connections are configured to mirror the relationship among the knowledge-base elements. Searching for information and unifying knowledge is then perfectly matched by the supporting VLSI.

Of the emerging parallel processing technologies, the Connection Machine System is probably one of the more esoteric. It usually requires special programming techniques that are, at best, little understood. And it makes little attempt to run standard software. But applications that can be adapted to its architecture have been shown to run orders of magnitude faster.

Origins Of The Machine

The Connection Machine is the brainchild of researcher Daniel Hillis, who developed the idea while at MIT. Prototypes of the Connection Machine have generated excitement among other researchers at MIT, who have been using it to experiment with parallel processing architectures.

An advantage of the machine in this type of research is the simplicity with which connection topologies can be set up. "The programmable interconnection greatly simplifies the problems of trying to get serial algorithms to emulate parallel architectures," said Charles Leiserson, a professor at MIT who works with VLSI architectures. "People around here are finding it indispensable for running simulations of architectures," Leiserson said.

Tomaso Poggio, an MIT researcher, claims that his system can reconstruct three-dimen-

sional scenes from two-dimensional information, using only about one quarter of the machine's 64,000 processors.

Other academics have also heaped praise on Hillis' invention. AI expert Marvin Minsky feels the system represents a fundamental break with computer-design tradition. Up to now, all computers have had basically the same architecture—one or a few large memory banks.

Challenging that conception, the Connection Machine links together thousands of millions of extremely small processors and memories. "It may take generations to unfold the implications of this new species of machine," Minsky said. He will be present at the formal unveiling this week.

But the question of how this machine will fare in an increasingly crowded parallel processing market now looms. One promising area that has been rapidly proven in research projects is VLSI design. "The machine's high degree of parallelism and general communications structure seems a natural fit to applications such as logic and circuit simulation, routing and placement, design rule checking, compaction and circuit extraction," Hillis said.

A processor can be allocated to each component in a circuit and the wiring pattern programmed into the Connection Machine's communication pathway to create a high degree of realism in a simulation.

—Additional reporting by
Chappell Brown

Thinking Machines Corp.
245 First St., Cambridge, MA
02142-1214; (617) 876-1111

Massively parallel processor introduced

Connection Machine concept represents a 'new technology'

By Eddy Goldberg

CAMBRIDGE, Mass. — The massively parallel processing Connection Machine introduced by Thinking Machines Corp. last week represents "a fundamentally new technology," according to company founder and President Sheryl Handler.

The Connection Machine incorporates 16,000 to 64,000 processors that reportedly process data at up to 1 billion instructions per second.

"I think it's a breakthrough, but I don't think it immediately makes other machines obsolete," said Vincent E. Giuliano, vice-president and chief scientist at Mirror Systems, Inc., which produces software systems for parent company Times Mirror Co.

"The important thing at this point is to prove the concept," said Martin Schultz, chairman of Yale University's computer science department.

"We have to find out if massive parallelism is the answer to our needs," he said. Schultz is looking to simulate three-dimensional physical experiments, among other tasks.

The company, founded in June 1983, has sold six units, which cost \$1 million in the 16,000-processor configuration and \$3 million for the full 64,000-processor version.

The first sale was to the Defense Advanced Research Projects Agency (DARPA), which sponsored the project to develop a general-purpose massively parallel computer.

"Each of the architectures we're sponsoring represents a major architectural advance," said Stephen Squires, assistant director of DARPA's Information Processing Technology Office. "This is pushing things in the direction of fine-grained, massively parallel, fully connected computing," he added.

Fine-grained processing assigns each data element of a task to an individual processor, allowing thousands of processors to work simultaneously, the company said. It works best where the number of data elements is large, ranging from 10,000

to one million.

The front end is a conventional computer, such as a Digital Equipment Corp. VAX or a Symbolics, Inc. 3600, which contains the programs. Any single-data-element instructions are executed directly by the front end. Instructions that call for operations on the whole data set at once are passed to the Connection Machine for execution.

For problems with more than 16,000 or 64,000 data elements, the processors, which each have 4,096 bits of memory, act as virtual processors. The Connection Machine can support up to one million virtual processors.

Its communications technology, called the router, has an overall capacity of 3 billion bit/sec. and allows any processor to establish a link to any other processor in a maximum of 12 dynamically configurable steps. Though its future usefulness is yet to be fully determined, four applications were demonstrated: document retrieval, fluid dynamics modeling, creating contour maps from aerial photographs and the design of very large-scale integrated circuits.